

This PDF includes a chapter from the following book:

The Cognitive Animal

Empirical and Theoretical Perspectives on Animal Cognition

© 2002 Massachusetts Institute of Technology

License Terms:

Made available under a Creative Commons
Attribution-NonCommercial-NoDerivatives 4.0 International Public License

<https://creativecommons.org/licenses/by-nc-nd/4.0/>

OA Funding Provided By:

The open access edition of this book was made possible by generous funding from Arcadia—a charitable fund of Lisbet Rausing and Peter Baldwin.

The title-level DOI for this work is:

[doi:10.7551/mitpress/1885.001.0001](https://doi.org/10.7551/mitpress/1885.001.0001)

20 Synthetic Ethology: A New Tool for Investigating Animal Cognition

Bruce MacLennan

Goals

Synthetic ethology is based on several methodological commitments. First, it is based on the conviction that the investigation of cognition should look at behavior and the mechanisms underlying that behavior in the agent's environment of evolutionary adaptiveness. Second, this investigation should extend over structural scales from the neurological mechanisms underlying behavior, through individual agents, to the behavior of populations, and over time scales ranging from neurological processes, through an agent's actions, to that of evolutionary processes. Obviously, such a wide range of scales is difficult to encompass in investigations of natural systems. Third is the observation that the discovery of deep scientific laws (especially quantitative ones) requires the sort of control of variables that can be achieved only in an artificial experimental setup.

Thus we are faced with conflicting demands. On the one hand, we need precise experimental control. On the other, ecological validity dictates that agents be studied in their environment of evolutionary adaptiveness, where there are innumerable variables that are not amenable to independent control. Synthetic ethology intends to reconcile these conflicting requirements by constructing a synthetic world in which the phenomena of interest may be investigated. Because the world is synthetic, it can be much simpler than the natural world and thereby permit more careful experimental control. However, although the world is synthetic and simple, it is nevertheless complete in that the agents exist, live, and evolve in it.

The original motivation for synthetic ethology came from one of the central problems in cognitive science: the nature of intentionality, the property that causes mental states to be about something. We felt that an understanding of intentionality would have to encompass both the

underlying mechanisms of intentional states and the social-evolutionary structures that lead to the creation of shared meaning. Our analysis of intentionality concluded that something is intrinsically meaningful to an agent when it is potentially relevant to that agent or to its group in its environment of evolutionary adaptedness (MacLennan 1992). Therefore intentionality must be studied in an evolutionary context.

We began our investigation with communication, since it involves both intentionality and shared meaning. We show in this essay how synthetic ethology permits the investigation of signals that are inherently meaningful to the signalers, as opposed to those to which we, as observers, attribute meaning.

Methods

The agents that populate our synthetic worlds can be modeled in many different ways; in particular there are a variety of ways of governing their behavior, including simulated neural networks and rule-based representations. In the experiments described here, an agent's behavior was controlled by a set of stimulus-response rules (64 rules in these experiments). These rules were determined by an agent's (simulated) genetic string, but they could be modified by a simple learning mechanism (described later).

Since our goal is to investigate the synthetic agents in their environment of evolutionary adaptedness, they must evolve. Our world includes a simplified form of simulated evolution, which proceeds as follows: Periodically two agents are chosen to breed, the probability of which is proportional to their "fitness" (as described later). The genetic strings of the two parents are mixed so that each of the offspring's genes comes randomly from either one or the other of the parents. In addition, there is a small probability of a gene being mutated. The result-

ing genetic string is used to create the stimulus-response rules for the single offspring, which is added to the population. In order to maintain a constant population size (100 in these experiments), one agent was chosen to “die” (i.e., to be removed from the population), the probability of dying being inversely related to “fitness.”

We illustrate here the sort of experimental control permitted by synthetic ethology. Because we have complete control over the experimental setup and the course of evolution, we may begin with genetically identical populations and observe their evolution under different experimental conditions. If something interesting is observed in the course of an experiment, we may rerun the exact course of the evolution of the population to that point, and then make additional observations or experimental interventions to investigate the phenomena. Finally, whenever any interesting phenomena are observed, there can be no fundamental mystery, for all the mechanisms are transparent. If some agent exhibits interesting behavior, its entire mechanism is available for investigation. There can be no “ghost in the machine.”

In synthetic ethology there is no requirement to model the natural world, as long as the synthetic world retains the essential characteristics of the natural world. That is, although determinate laws govern the evolution of our experimental populations, we are able to decide our world’s “physical laws,” which determine whether an agent “lives” or “dies,” and which select agents for reproduction. The goal, of course, is to create synthetic worlds that are like the natural world in relevant ways, but are much simpler to study. The following experiment illustrates what can be accomplished.

Demonstrating the Evolution of Communication

Methods

Our first series of experiments investigated whether it was even possible for genuine, mean-

ingful communication to evolve in an artificial system. We decided to construct the simplest possible system that could be expected to lead to real communication.

Although there are many purposes for which an agent might be expected to communicate, we decided to focus on cooperation. Our reasoning was that communication could be expected to evolve in a context in which some agents have information that other agents could use to facilitate cooperation. Thus we gave each agent a local environment that could be sensed by that agent but by no other. It can be thought of as the situation in an animal’s immediate vicinity, but to keep the model as simple as possible, we limited the local environments to be in a small number of discrete states (eight in these experiments).

To make the state of one agent *B*’s local environment relevant to another agent *A*, we arranged that they could cooperate only if *A* performed an action suitable for *B*’s environment. To make this cooperation as simple as possible, we made our agents capable of producing an action from the same set as the local environment states. Thus *A* could cooperate with *B* only by producing the same item that was in *B*’s local environment, which *A* could not sense directly.

To select for cooperation, we simply measured the number of times, over a specified period, that each agent was involved in successful cooperation. The probability of an agent reproducing was made proportional to this rate of cooperation, and its probability of dying was inversely related to the rate in a simple way. Thus we placed selective pressure on cooperation, but not directly on communication; indeed, limited cooperation can be achieved by random action (which has a 1:8 chance of succeeding).

Our experiments implemented only micro-evolution, so our agents were unable to evolve new sensor or effector organs. We gave our agents organs that might be used for communication, but we did not construct the agents to use them in this or any other way.

Again, simplicity was our principal aim. We equipped our synthetic world with a simple global environment, shared by all the agents, which could be in one of a few discrete states (eight in these experiments). The agents had the physical capability of sensing and modifying this global environment. Specifically, the state of the global environment is part of the stimulus to which an agent reacts, and the response can be either a new state for the global environment or an attempt to cooperate.

To test the potential effects of communication on cooperative behavior, we implemented a mechanism for making communication impossible. When communication was being suppressed, we periodically randomized the state of the global environment. This allowed us to measure the effect of apparent communication on the fitness (rate of cooperation) of the population, since genuine communication is defined in terms of its effect on the fitness of the communicators (Burghardt 1970).

We also investigated a very simple form of single-case learning, which could be enabled or disabled. When it was enabled, learning took place when an agent attempted to cooperate but failed. In this case, the stimulus-response rule used was changed to what would have been correct in this situation (although there was no guarantee that it would be the correct response in the future).

We initialized our population with 100 individuals containing random genetic strings. Thus the stimulus-response rules governing their behavior, which were determined by their genomes, were also initially random.

Results

To be able to measure the effect of communication on the fitness of a population, we quantified the fitness by the number of successful cooperations per unit of time, which we called the “degree of coordination” of the population. (The unit of time was a “breeding cycle,” in which one individual died and one was born.) Because there

was considerable random variation in the degree of coordination, the time series was smoothed by a moving average. Linear regression was used to establish the rate at which the degree of coordination (fitness) increased or decreased. Details can be found in MacLennan (1990, 1992) and MacLennan and Burghardt (1993).

The baseline for comparison was determined by suppressing all possible communication, as previously described. In this case the degree of coordination stayed near 6.25 cooperations per unit of time, the level the analysis predicted would occur in the absence of communication. Linear regression showed a slight upward trend in the degree of coordination (the reason for which is discussed in the papers cited).

On the other hand, when communication was not suppressed, we found that the degree of coordination increased 26 times faster than when communication was suppressed. Over an interval of 5000 breeding cycles, the degree of coordination reached 10.28 cooperations per unit of time, which is 60 percent higher than the 6.25 achieved when communication was suppressed (figure 20.1). When the agents were permitted to learn from their mistakes, fitness increased at 3.82 times the rate found when learning was disabled, and at approximately 100 times the rate that occurred when communication was suppressed.

As would be expected for experiments of this kind, there is considerable experimental variation from run to run. Nevertheless, the results we have described are typical over more than a hundred experiments. Therefore we can conclude that genuine, meaningful communication is taking place, for it is significantly enhancing the fitness of the population. Furthermore, since communication evolves in our population when it is not suppressed, we may investigate genuine communication in its environment of evolutionary adaptedness.

Since it is genuine communication, the signals passed among the agents are meaningful to them, but not necessarily to us as observers. That is, we have a situation opposite from that of artificial intelligence, in which the computer




Figure 20.1
Increase in degree of coordination that is due to evolution of communication.

manipulates symbols that are meaningful to us but meaningless to it (or, more precisely, they have only a derived meaning that is dependent on the meaning we attribute). Here we are in the same situation as in natural ethology; we are faced with apparently meaningful communication and must discover its meaning *for the communicators*.

Even in these simple experiments, signals and their interpretation are complex functions of the total situation. The signal emitted by an agent may depend on both its local environment and the shared global environment. Furthermore, an agent's interpretation (use) of a signal may be (and typically is) influenced by its own local environment.

Nevertheless, we would expect that over time a simple meaning would emerge for the signals; that is, that there would be a one-to-one correspondence between signals and local environ-

ment states. To determine if this was occurring, we compiled a co-occurrence table, which recorded the number of times particular pairings of a signal (global environment state) and meaning (local environment state) occurred in successful cooperations.

If no communication is taking place, one would expect all signal-meaning combinations to be about equally likely, and that is what we found when communication was suppressed, and at the beginning of the simulations when it was not. However, when communication was not suppressed, the co-occurrence tables became more structured as the "language" self-organized.

We quantified the organization of the co-occurrence tables in a number of different ways, including entropy, a measure of disorder (lower numbers represent greater organization). With our experimental design, the maximum entropy,

when all signal–meaning pairs are equally likely, is 6 bits, but when there is a one-to-one symbol–meaning correspondence, the entropy is 3 bits. When communication was suppressed, we observed an entropy of 4.95 bits, which shows that it was not completely disordered. However, when communication was not suppressed, the entropy decreased (after 5000 breeding cycles) to 3.87, representing a much higher degree of organization.

Visual inspection of the evolved co-occurrence tables showed a number of cases in which, almost always, a particular signal corresponded to a particular meaning and vice versa. However, we also observed cases of ambiguity, in which a signal was more or less equally likely to correspond to two or more meanings, and cases of synonymy, in which two or more signals were about equally likely to correspond to a particular meaning. These cases could result from individual agents using ambiguous or synonymous symbols, or from two or more competing “dialects” in the population, but Noble and Cliff (1996) have evidence supporting the former hypothesis.

The observations described here can be called “behavioral” and are analogous to those made in natural ethology. However, synthetic ethology affords additional possibilities, for the structure of the agents is completely transparent. At any time we may “dissect” the agents and analyze their behavioral programs (see MacLennan 1990 for examples). Thus we may relate the mechanisms of behavior to their manifestation in the population.

Brief Overview of Other Experiments

We have been interested in whether a population would evolve to use sequences of symbols for communication if there was a selective advantage in doing so. To explore this possibility, my students and I have conducted a number of experiments similar to those already described; details may be found in MacLennan (2001) and in the references cited there. In these experiments

the agents evolved the ability to communicate with pairs of symbols displaying a very rudimentary “syntax,” but the results have been less interesting than we expected. One explanation may be that the very simple behavioral model we used (finite-state machines) is too weak for the sequential perception and control required for more complex communication. Animals, in contrast, have rich, highly structured perceptual-motor systems, which evolution can recruit for communication. Future experiments might need to use more complex models of agents, as well as a more structured environment about which they might communicate.

Discussion

Of necessity, our discussion of related and future work and its implications must be brief. Noble and Cliff (1996) have replicated our earliest studies and extended them in a number of informative ways. A somewhat different approach can be found in Werner and Dyer (1992), who demonstrated the evolution of communication by making it necessary for effective reproduction. Steels (1997a, b) has conducted fundamental studies on the emergence of meaningful symbols.

In discussing related work, it may be worth making a few remarks about the connection between synthetic ethology and a related discipline, artificial life. First, it must be stressed that there is substantial overlap between them, so that the difference is at most one of emphasis.

Artificial life studies artificial systems that are significantly “lively” in some sense. Some investigators are attempting to create systems that are literally alive, while others are content with systems that faithfully imitate life. Synthetic ethology differs from this discipline in that the agents need not be alive in either of these senses, although they may be. Certainly we make no claim that the agents described in this essay are alive in any literal sense.

Current experiments in synthetic ethology are too simple to exhibit psychological states, but future ones may be able to do so; even the current experiments exhibit genuine intentionality. Synthetic ethology indicates how psychological states may be made accessible to scientific investigation.

We have claimed that our agents (although they are not conscious or even alive) exhibit genuine intentionality. The point is certainly arguable and depends on our analysis of intentionality. Nevertheless, all subtleties aside, we claim that the signals are inherently meaningful to the agents because the agents' continued persistence as organized systems depends on their use of the signals.

Are these synthetic worlds and agents so alien that results will not be seen as relevant to nature? In particular, we have argued that we can use abstract, ad hoc selection rules (since the "laws of physics" are under our control), but it can be objected that selection should be more naturalistic (e.g., Werner and Dyer 1992). Certainly this is an important issue, and in the long run we want to explore ever richer synthetic worlds, but introducing gratuitous complexity would defeat the goals of synthetic ethology.

One of the advantages of synthetic ethology is that we can make our worlds as simple as possible, as long as they include the phenomena of interest. On the other hand, we must construct these worlds from scratch; they are not given to us. This becomes a challenge as we begin to investigate phenomena that require larger populations of more complex agents acting in more complex environments. The simulation of such worlds requires ever more powerful computers. Therefore synthetic ethologists must strike a delicate balance between the sophistication of the synthetic world and the resources required to implement it. Indeed, as we move in the direction of greater complexity, synthetic ethology will face some of the same problems as natural ethology. Nevertheless, by affording greater control and an

alternative to natural life, it will remain a worthwhile approach.

References

- Burghardt, G. M. (1970). Defining communication. In *Communication by Chemical Symbols*, J. W. Johnston, Jr., D. G. Moulton, and A. Turk, eds., pp. 5–18. New York: Appleton-Century-Crofts.
- MacLennan, B. J. (1990). *Evolution of Communication in a Population of Simple Machines*. Technical Report No. CS-90-99. Knoxville, Tenn.: Computer Science Department, University of Tennessee.
- MacLennan, B. J. (1992). Synthetic ethology: An approach to the study of communication. In *Artificial Life II: The Second Workshop on the Synthesis and Simulation of Living Systems*, C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, eds., pp. 631–658. Redwood City, Calif.: Addison Wesley.
- MacLennan, B. J. (2001). The emergence of communication through synthetic evolution. In *Advances in Evolutionary Synthesis of Neural Systems*, V. Honavar, M. Patel, and K. Balakrishnan, eds., pp. 65–90. Cambridge, Mass.: MIT Press.
- MacLennan, B. J. and Burghardt, G. M. (1993). Synthetic ethology and the evolution of cooperative communication. *Adaptive Behavior* 2: 161–188.
- Noble, J. and Cliff, D. (1996). On simulating the evolution of communication. In *From Animals to Animals 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, P. Maes, M. Mataric, J.-A. Meyer, J. Pollack, and S. W. Wilson, eds., pp. 608–617. Cambridge, Mass.: MIT Press.
- Steels, L. (1997a). Constructing and sharing perceptual distinctions. In *Proceedings of the European Conference on Machine Learning*, M. van Someran and G. Widmer, eds., pp. 4–13. Berlin: Springer-Verlag.
- Steels, L. (1997b). The synthetic modeling of language origins. *Evolution of Communication* 1: 1–34.
- Werner, G. M. and Dyer, M. G. (1992). Evolution of communication in artificial organisms. In *Artificial Life II: The Second Workshop on the Synthesis and Simulation of Living Systems*, C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, eds., pp. 659–687. Redwood City, Calif.: Addison Wesley.